

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-03-

0173

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering the required data, reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Project, Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE	3. REPORT TYPE AND DATES COVERED 1 DEC 99 - 30 NOV 02	
4. TITLE AND SUBTITLE HANDLING CONTRADICTIONARY DATA WITH METREASONING			5. FUNDING NUMBERS F49620-00-1-0068	
6. AUTHOR(S) DON PERLIS				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) UNIVERSITY OF MARYLAND COLLEGE PARK, MD 20742			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 4015 Wilson Blvd, Room 713 Arlington, VA 22203-1954			10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-00-1-0068	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Our work has involved extending the meta-reasoning capabilities of the Active Logic inference engine (ALMA) to be able to represent a variety of types of contradiction and positive and negative introspection. This involved both theoretical and implementational effort. In particular, we developed techniques for representing, recognizing, and repairing mistakes that arise in the form of contradictions in the evolving belief set of an agent. This had two aspects: general (commonsense) reasoning, and reasoning attendant to natural-language dialogue. This work led to a number of other publications (see below), including a PhD dissertation completed in 2001.				
14. SUBJECT TERMS			15. NUMBER OF PAGES 8	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

20030520 113

Herklitz

FINAL REPORT

Project: Handling Contradictory Data With Metareasoning

PI: Donald Perlis (and Michael Anderson, co-PI)
University of Maryland
College Park, MD 20742

Agreement # F496200010068

OBJECTIVES:

As our principal research objective, we intend to isolate, formalize, implement, and test effective means for detecting and repairing contradictory information within intelligent reasoning systems in real-time. The active logic approach, developed in prior work, provides a framework in which contradictions in the system's knowledge base will not inhibit the system from performing meaningful inference. We intend to augment the basic framework with specific techniques for successfully recognizing and adjudicating these contradictions as they arise. We will implement and test the contradiction handling mechanisms within an automated reasoning system, and finally we will apply these means to solving problems in two test domains: real-time planning and execution, and human-computer dialogue interaction.

STATUS OF EFFORT:

See Accomplishments, below. The grant period ended December, 2002.

SUMMARY:

Our work has involved extending the meta-reasoning capabilities of the Active Logic inference engine (ALMA) to be able to represent a variety of types of contradiction and positive and negative introspection. This involved both theoretical and implementational effort. In particular, we developed techniques for representing, recognizing, and repairing mistakes that arise in the form of contradictions in the evolving belief set of an agent. This had two aspects: general (commonsense) reasoning, and reasoning attendant to natural-language dialogue. This work led to a number of other publications (see below), including a PhD dissertation completed in 2001.

ACCOMPLISHMENTS:

1. We extended our formal techniques for representing the reasoning of multiple agents, to allow for intentions. In this work, we viewed the reasoning of agents as ongoing processes rather than as fixed sets of conclusions. Our approach utilizes a strongly sorted calculus, distinguishing the application language, time, and various syntactic sorts. We have established soundness and completeness results corresponding to various families of agents. This allows for useful and intuitively natural characterizations of such agents' reasoning abilities. We discuss and contrast consistency issues as in the work of Montague and Thomason. We also showed how to represent the concept of focus of attention in this framework. This work was published in *Autonomous Agents and Multi-Agent Systems*.

The major examples we use to illustrate the use of this framework are (i) preparing to make a phone call, and (ii) planning to get to the airport. These examples can also be viewed as deadline problems, which makes them especially relevant to real-world settings.

Suppose that an agent B is in Israel and wants to phone Jack. B knows Jack's local phone number, but not Jack's country code and region (area) code. B also knows that there is an agent C that knows country codes and an agent A that knows area codes. Hence B asks agent C for Jack's country code and agent A for Jack's area code.

Agent C knows the country codes of various people and has general knowledge about country codes. In particular, agent C knows that Jack and Jill live in the same country, that Jill lives in the USA, and that the USA country code is 1.

We showed that in our formalism, we can establish conditions under which it is possible to prove that agent B will, be able to ascertain that, with the help of agents A and C, it is possible to determine Jack's phone number. While this result seems intuitively obvious, it in fact depends on knowledge, by B, of rather complex interactions between the three agents and especially their states of knowledge. This is particularly so when the inferences are defaults, ie, when they are based on plausible but possibly false assumptions.

2. We have prepared a draft ALMA User's Manual.

3. Summary of Dissertation.

Khemdut Purang, "Systems that Detect and Repair their own Mistakes", University of Maryland 2001.

Making mistakes is an inescapable aspect of everyday life. We constantly make mistakes, recognize them and try to correct them. Mistakes are inevitable because of the incompleteness of our knowledge of the world, its inherent uncertainty and its being in a constant state of change. We can never know for sure that what we know is true and the actions that we take based on these beliefs can therefore be misguided. Sooner or later we act based on some false belief or the world changes in an unexpected way and we fail to achieve our goal. But the fact that we can recognize and repair these errors mitigates their effects. Software systems face the same problems. The difference is that they do not usually have as robust a capability as we have to detect and respond to their mistakes. This is part of what makes them seem brittle and user-unfriendly. This problem is not likely to get any better as the systems exhibit more complex behaviors in more realistic domains. Our work begins to address that problem by focusing on the computational capabilities required of

software systems for them to be able to autonomously recognize and respond to their own mistakes. We study in particular, mistaken beliefs, intentions and actions in agents that have some goals to achieve. Intuitively enough, the basic capabilities required are an ability to inspect their past behavior and computations and the past states of the world and to use that to determine their future behavior. These abilities are not typically available in software systems. We have implemented a general logical framework in which one can specify the behavior of an agent that supports this kind of representation and computation. We have implemented agents that detect and respond appropriately to their mistakes in some aspects of language processing. We have also implemented a system that handles its mistaken beliefs in any domain that can be described using the language of non-monotonic logic. This system was tested on a test suite that we compiled from examples of non-monotonic reasoning in the literature. We finally provide a design of the representations and algorithms for handling mistakes in an agent that acts in the world and has mistaken beliefs, intentions and actions. Implementing such an agent is the next step in this work.

4. We designed and implemented extensions of the existing active logic engine and dialogue manager, to be able to engage in a variety of action directive subdialogues, which include:

- (a) a request for action performance by the user
- (b) performance of an action by the system
- (c) feedback ("followup") from the user (which might be positive indicating acceptance of the action, or negative, possibly coupled with a replacement action).

We also designed a more powerful architecture that will allow us to process new words and phrases and implicit and explicit quotation, in order to adapt to a variety of users and contexts. As a result, users will be able to instruct the system in exchanges such as:

- (d) I use "transportation" to refer to any vehicle, not just a car.
- (e) Peking is the old name for Beijing.

5. We installed the JESS programming system from SANDIA National Laboratories, and have begun work on a new version of ALMA in JESS, which will be web-accessible.

6. We performed a new study, combining metareasoning and learning modules, so that as learning proceeds, the metareasoner can monitor its progress and offer guidance. We implemented a simple version of this, for a simulated robot navigating a maze, with overall speed as the target variable to maximize. Whenever the robot bumps a wall, the metareasoner issues a command to the learner, to retrain. After many trials, the robot's navigational behavior improves so that it no longer bumps walls. While learned behavior is a well-studied topic, its guidance by a reasoning module is new, and our planned further work should show that it has considerable power over traditional methods. This work was done using the JESS system for the reasoning component.

PERSONNEL SUPPORTED

D Perlis (PI)
M Anderson (co-PI)
K Purang (GRA)
Y Chong (GRA)
Y Okamoto (GRA)
D Josyula (GRA)
A Rogriguez (GRA)
G Lapizco (GRA)

TRANSITIONS (to DoD, Government, Industry):

None.

PUBLICATIONS:

1. A logic for characterizing multiple bounded agents (in Autonomous Agents and Multi-Agent Systems, 3, 351--387, December 2000).
2. Representations of Dialogue State for Domain and Task Independent Meta-Dialogue. David Traum, Carl Andersen, Yuan Chong, Darsana Josyula, Yoshi Okamoto, Khemdut Purang, Michael O'Donovan-Anderson and Don Perlis. Electronic Transactions on AI, to appear, 2002.
3. Alma: an interactive Active Logic Machine. Khemdut Purang. IJCAI '01 Workshop on Effective Interactive AI Resources.
4. Zi at the Interface and Beyond. Yoshi Okamoto. Presented at the LOT-NWCL Summer School. Utrecht, the Netherlands: Universiteit Utrecht, 2001.
5. Seven Days in the Life of a Robotic Agent. Waiyian Chong, Mike

O'Donovan-Anderson, Yoshi Okamoto and Don Perlis.
First GSFC/JPL Workshop on Radical Agent Concepts,
2001, NASA Goddard Space Flight Center, Greenbelt, MD, USA

6. Handling Uncertainty with Active Logic, Manjit Bhatia, Paul Chi, Waiyian Chong, Darsana P. Josyula, Yoshi Okamoto, Don Perlis and Khemdut Purang.
2001 AAAI Fall Symposium Series in North Falmouth, MA.

7. Agents that recognize and repair their own mistakes. PhD
dissertation. Khemdut Purang. August 2001.

8. Symbol Systems. With Michael L. Anderson.
Encyclopedia of Cognitive Science. 2002.

9. A Logic-Based Model of Intentions for Multi-Agent Subcontracting.
AAAI-2002 With John Grant and Sarit Kraus.

10. Invited one-hour talk and paper on self-reference, to be presented
at PhiLog, Copenhagen, October 2002.

11. The Use-Mention Distinction and its importance to HCI. Michael
L. Anderson, Yoshi Okamoto, Darsana Josyula, and Don Perlis. Proceedings of
the Sixth Workshop on the Semantics and Pragmatics of
Dialog. 2002.

12. Time-Situated Agency: Active Logic and Intention Formation.
To appear, Cognitive Agents Workshop, 2002. With M. Anderson,
D. Josyula, Y. Okamoto.

13. Meta-reasoning for Conversational Adequacy. Michael L. Anderson
Invited talk, Deutsches Forschungszentrum fr Knstliche Intelligenz. 2002.

INTERACTIONS:

1. invited talk at New Mexico University

2. contact and discussions with ARL (Dr. John Gurney)
3. invited talk at PhiLog (Logic Colloquium in Denmark)
4. invited talk at DFKI.

Inventions: none

Awards: none